

Decisions from Experience:

How Groups and Individuals Adapt to Change

Tomás Lejarraga

Max Planck Institute for Human Development, Berlin, Germany

José Lejarraga

IE Business School, IE University, Madrid, Spain

Cleotilde Gonzalez

Carnegie Mellon University, Pittsburgh, USA

Author Note

Tomás Lejarraga, Center for Adaptive Rationality, Max Planck Institute for Human Development; José Lejarraga, IE Business School, IE University; Cleotilde Gonzalez, Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University.

This research was supported in part by a grant from the Defense Threat Reduction Agency (DTRA) (HDTRA1-09-1-0053) and by National Science Foundation award number 1154012 to Cleotilde Gonzalez. We are grateful for discussions with and comments from Tara Wernsing, Taya Cohen, Tilmann Betsch, and Juliane Kämmer. We thank Hau-yu Wong for research assistance and Susannah Goss for editing the manuscript.

Correspondence concerning this article should be addressed to Tomás Lejarraga, Center for Adaptive Rationality, Max Planck Institute for Human Development, Lentzealle 94, 14195, Berlin, Germany. E-mail: lejarraga@mpib-berlin.mpg.de

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>2014</b>		2. REPORT TYPE		3. DATES COVERED <b>00-00-2014 to 00-00-2014</b>	
4. TITLE AND SUBTITLE <b>Decisions from Experience: How Groups and Individuals Adapt to Change</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Carnegie Mellon University,Dynamic Decision Making Laboratory,Department of Social and Decision Sciences,Pittsburgh,PA,15213</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>					
13. SUPPLEMENTARY NOTES <b>In press. Memory &amp; Cognition.</b>					
14. ABSTRACT <b>Whether groups make better judgments and decisions than individuals has been studied extensively, but most of this research has focused on static tasks. How do groups and individuals compare in settings where the decision environment changes unexpectedly and without notification? This article examines group and individual behavior in decisions from experience where the underlying probabilities change after some trials. Consistent with the previous literature, the results showed that groups performed better than the average individual while the decision task was stable. However, group performance was no longer superior after a change in the decision environment. Group performance was closer to the benchmark of Bayesian updating, which assumed perfect memory. Findings suggest that groups did not adopt decision routines that might have delayed their adaption to change in the environment. Rather, they seem to have coordinated their responses, which led them to behave as if they had better memory and subsequently delayed adaptation.</b>					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES <b>38</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

### Abstract

Whether groups make better judgments and decisions than individuals has been studied extensively, but most of this research has focused on static tasks. How do groups and individuals compare in settings where the decision environment changes unexpectedly and without notification? This article examines group and individual behavior in decisions from experience where the underlying probabilities change after some trials. Consistent with the previous literature, the results showed that groups performed better than the average individual while the decision task was stable. However, group performance was no longer superior after a change in the decision environment. Group performance was closer to the benchmark of Bayesian updating, which assumed perfect memory. Findings suggest that groups did not adopt decision routines that might have delayed their adaption to change in the environment. Rather, they seem to have coordinated their responses, which led them to behave as if they had better memory and subsequently delayed adaptation.

*Keywords:* group versus individual decisions, decisions from experience, uncertainty, changing environments, adaptation, instance-based learning

### Decisions from Experience:

#### How Groups and Individuals Adapt to Change

Every morning, Ferran Adrià receives fresh products from his suppliers on Spain's Costa Brava. High quality raw ingredients are crucial to the success of El Bulli, once ranked as the world's best restaurant. Adrià monitors the quality of his suppliers on a daily basis, and makes decisions about which suppliers to keep and which to change according to the quality of the goods they provide. Although the quality of the products fluctuates—after all, it depends on random factors like the weather—a sustained drop in quality must be detected and acted upon quickly. Given the importance of this task for El Bulli, would the restaurant benefit from having a group, instead of a single individual, involved in this decision? Perhaps. Only a few kilometers from El Bulli, Joan, Josep, and Jordi Roca manage El Celler de Can Roca, currently considered the best restaurant in the world (“The World's 50 Best Restaurants,” 2013).

Whether groups make better judgments and decisions than individuals has been studied extensively. The common observation is that groups perform better than the average individual in a variety of intellectual tasks (Hill, 1982). This consistent result is believed to occur because groups have higher information-processing capacity: groups are able to gather more information (Cohen & Thompson, 2011), make fewer mistakes in processing it (Charness & Sutter, 2012), and use it more consistently than individuals (Hinsz, Tindale, & Vollrath, 1997). Although past research has made substantial progress in understanding how group and individual behavior compares, most of this research has focused on behavior in static environments, and little is known about their relative performance in settings where the decision environment changes unexpectedly and without notification.

Coordinating different opinions within a group is costly and often detrimental to group performance (Steiner, 1972). When decisions are recurrent, however, group members are able to

develop routines that allow them to evade a costly coordinated response (Gersick & Hackman, 1990). A routine is a learned behavioral solution that comes to mind first when a familiar decision situation is encountered (Betsch, 2005). Routines can, however, become detrimental if the decision environment changes, because new decision situations may need new solutions. A routine developed under particular conditions may not produce an adequate response under different conditions (Betsch, Haberstroh, & Höhle, 2002). For example, a restaurateur who chooses his suppliers based on long-term reputation (i.e., a routine) may overlook a new, high-quality supplier.

Routines can be particularly detrimental in environments whose structure is not directly observable, and where decision makers are only able to learn about the properties of their options by experiencing the outcomes of their choices. Such environments have been the focus of recent research and are referred to as *decisions from experience* (Hertwig, Barron, Weber, & Erev, 2004). How do groups and individuals compare when coping with change in decisions from experience?

### **Decisions from Experience: A Suitable Paradigm for Studying Adaptation to Change**

In everyday life, people often lack explicit information when making decisions; instead, they have to rely on past experience and adapt their decisions accordingly (Gonzalez, Lerch, & Lebiere, 2003; Hertwig & Erev, 2009). Managers like Adrià learn about the quality of their suppliers by observing recurrent feedback. They decide whether or not to continue a relationship with a given supplier on a daily basis. Such decisions from experience have been studied experimentally since the 1960's (Katz, 1964; Myers & Sadler, 1960). In a common experimental paradigm, respondents see two buttons on a computer screen, each one representing a payoff distribution unknown to participants (e.g., 4 with 80% chance or 0 otherwise). Respondents

choose between these two distributions repeatedly for a fixed number of trials (e.g., 100 trials), receiving immediate feedback on the outcomes of their choices. They thus gradually learn about the payoffs of the two options from outcome feedback.

In most studies of decisions from experience, outcome distributions are constant across all trials (e.g., the probability of 4 is constant at 80% in all trials). However, Rakow and Miler (2009) used a variant of this paradigm to study how individuals adapt to unexpected changes in the probabilities of outcomes. In their experimental design, the outcome probabilities of one of the two options changed gradually across trials, while those of the other option remained constant. This design made one option objectively better than the other in early trials, but worse in later trials. It allowed the authors to study whether and how respondents adapted their choices as one option changed from better to worse. One key aspect of this experimental task is that the change in the outcome probabilities was not announced, neither was it immediately obvious from observing the outcomes of one's choices. This problem thereby challenged participants to detect a change in probabilities when those probabilities could only be gauged by observing random draws from the outcome distribution.

The experimental design introduced by Rakow and Miler (2009) is thus a functional abstraction of the decision situation described in the introduction. We used their approach to study how groups and individuals compare when adapting their choices to an unexpected change in the decision environment. Throughout this article, we refer to adaptation as the ability to identify and select an option that has changed from worse to better relative to another option across trials, and to abandon an option that has changed from better to worse.<sup>1</sup>

### **Adapting to Change in Decisions from Experience**

---

<sup>1</sup> This view of adaptation differs from the biological approach, which describes how organisms evolve by means of natural selection.

The relative performance of groups and individuals in adapting to change in decisions from experience is an unexplored area in the literature on judgment and decision making. Rakow and Miler (2009) explored individual adaptation, and found that individuals adapted “slowly” to payoff changes. Moreover, they showed that participants who were made aware of past outcomes (through summary descriptions) made better choices than did individuals without memory support when the environment was stable, but most of the times they adapted less well after the environment changed. Longer lasting memory is desirable in stable environments, where the best predictor of future outcomes is a large sample of past observations, but it may be undesirable in an environment that has changed, where recent outcomes are more representative of future outcomes than are older ones. Recency, then, may have helped individuals adapt to changes.

Although recency effects are a common finding in the context of decisions from experience (Barron & Erev, 2003; Erev, Ert, & Yechiam, 2008; Erev et al., 2010; Hertwig et al., 2004; Lejarraga, Dutt, & Gonzalez, 2012), they have been observed only in individuals; it remains uncertain how groups use past outcomes in their decisions. The memory advantage of groups is well established in the literature (Betts & Hinsz, 2010). It has been attributed partly to better error correction (e.g., Hinsz, 1990; Vollrath, Sheppard, Hinsz, & Davis, 1989), but evidence also suggests that groups rely more heavily than individuals on memorization (Olsson, Juslin, & Olsson, 2006). This higher reliance on better memory suggests that groups may be at a disadvantage relative to individuals when making decisions from experience in a changing environment.

Groups tend to accentuate the behaviors of their members, and if members are likely to process information in a specific way, a group formed by those members is therefore even more likely to process information in that way (Hinsz et al., 1997). Reimer, Bornstein, and Opwis

(2005) examined this “accentuation” hypothesis in a problem-solving task with an explicit solution: the Tower of Hanoi game. They explored how dyad and individual behavior compared in the context of moving from one set-up to another. To this end, the authors taught their participants to solve the problem following one of two procedures. One procedure was goal-driven and consisted in dividing the problem into subproblems to be solved sequentially. The other was stimulus-driven and based on a simple rule of movements. Although both procedures guaranteed an identical solution in the first set-up, only the goal-driven procedure could be directly applied to a different set-up. The authors showed that, when participants learned the goal-driven procedure—which was not easily transferrable—both dyads and individuals took time to adapt, but dyads took longer. We might therefore expect groups to take longer to adapt to changes in the environment in decisions from experience.

However, Kämmer, Gaissmaier, and Czienskowski (2013) found that the relative adaptation of groups and individuals is task-dependent. The authors used a multi-attribute choice task in which respondents selected one of two hypothetical oil-drilling sites on the basis of various cues. They created a decision task with two environments. In the one environment, relying on a single cue led to more optimal choices; in the other, integrating all cues was optimal. Participants made choices in one environment and then in the other, in a randomized order. When participants began the task in the environment that favored a single-cue strategy, groups and individuals adapted their use of strategies well and to the same extent. However, groups adapted better than individuals when they began the task in the environment that favored a multiple-cue strategy. The authors argued that the groups’ better performance could be attributable to their superior capacity as information processors.

### **Routines and group coordination**



Routines are common in everyday decision making, as they save the time and effort of constantly considering the best possible course of action (Betsch, Fiedler, & Brinkmann, 1998; Betsch, Haberstroh, Glöckner, Haar, & Fiedler, 2001; Betsch et al., 2002; Schneider & Shiffrin, 1997). These learned solutions can be established at the level of options (Betsch et al., 2001) or of decision strategies (Bröder & Schiffer, 2006), depending on the characteristics of the situation (Bröder et al, 2013). For example, an individual or group may learn to choose a particular option because it results in an attractive outcome most of the time (routine option), but they may also learn a specific decision rule, such as staying with an option if the most recent outcome was attractive and otherwise switching (routine strategy). Although routine options and strategies may be adopted by both individuals and groups, the latter are more likely to adopt routines of any type, because without a routine, group members would have to coordinate a decision in a case-by-case manner—something unnecessary for individuals deciding alone. Routines can therefore be expected to make groups more resistant than individuals to updating their choices once a change has been detected. Indeed, diversity within a group has been found to increase the cost of reaching consensus and often leads to choice refusal (Nijstad & Kaps, 2008), a behavior that is consistent with the adoption of routines.

### **A benchmark model**

We use the task proposed by Rakow and Miler (2009) to examine how participants adapt to changes. In this task, participants are not told that probabilities will change, but they are made aware that outcome probabilities *may* change (see instructions below). Participants remain uncertain whether probabilities will actually change, where they could change (which option), and when they could change (which trial). In this context, we assume that the problem has not

changed until sufficient evidence accumulates to indicate otherwise.<sup>2</sup> We thus used a Bayesian updating model as a benchmark to compare the behavior of individuals and groups. According to this Bayesian updating model, the value  $E$  of an option  $j$  at trial  $t$  is:

$$E_j(t) = \frac{\phi + N_j(t-1)}{\phi + N_j(t)} \cdot E_j(t-1) + \frac{1}{\phi + N_j(t)} \cdot u(t)$$

where  $N_j(t)$  is the number of outcomes observed in option  $j$  up to and including trial  $t$ ;  $u(t)$  is the value of the outcome in trial  $t$ ; and  $\phi$  is the weight given to initial expectations. The initial expectations are beliefs about  $E_j$  in the very first trial. However, by setting  $\phi = 0$ , as we do here, the model ignores initial expectations and makes  $E_j(t)$  fully dependent on the observed outcomes. The model, then, assumes perfect memory and weighs all previous outcomes equally: It becomes a running mean (Biele, Rieskamp, & Gonzalez, 2009; Yechiam & Busemeyer, 2005). The model chooses the option with the higher value  $E$ , and chooses randomly (with .5 chance) if both options have the same  $E$ .<sup>3</sup>

Using a Bayesian updating model allows us to identify different processes of adaptation. Participants may switch to the better option after the environment has changed, but in line with the indication of Bayesian updating (what Betsch, et al., in press, refer to as *routine effects*). Alternatively, they may deviate from Bayesian updating in two ways: They may delay adaptation

---

<sup>2</sup> The same assumption was used by Betsch, Lindow, Engel, Ulshöfer, and Kleber (in press) in a similar task where change was not explicitly announced.

<sup>3</sup> This model is equivalent to a model that calculates expected values in each period using observed outcomes and probabilities derived from Bayes Rule,  $p(g|o) = \frac{p(o|g) \cdot p(g)}{p(o)}$ , where  $p(g|o)$  is the conditional probability of a gain given the participant chose the option  $o$  being evaluated;  $p(g)$  is the count of gains (in both options) divided by the number of choices made;  $p(o)$  is the count of times the participant has observed outcomes from  $o$  divided by the number of choices made; and  $p(o|g)$  is the total number of gains while choosing  $o$  divided by the count of gains (in both options). Then, the probability of gain given the participant chooses  $o$  is  $p(g|o) = \frac{\text{gains in } o}{\text{number of choices}}$ .

beyond Bayesian updating (what Betsch, et al., in press, refer to as *routine biases*) or they may adapt to a change before this benchmark.

According to this Bayesian updating model, the point to change from one option to the other is not when the change in the environment occurs, but later, when sufficient evidence of the change has been compiled. If groups have a memory advantage, as suggested by previous research, they can be expected to behave more in accordance with Bayesian updating than individuals, because Bayesian updating assumes that all previous outcomes are considered before the current value of an option is updated. At the same time, if individual decision-making behavior shows recency effects, as is commonly observed in the literature, individuals can be expected to deviate from Bayesian updating, adapting more quickly than the benchmark.

#### **What implications do gradual versus sudden changes have for adaptation?**

Differences in reliance on memory may lead to differences in adaptation when changes occur suddenly versus gradually. A gradual change in the probabilities of outcomes, as explored by Rakow and Miler (2009), makes past experiences decreasingly representative of the current situation. That is, past experiences have some informational value for decisions in the near future, and groups—with their better memory—may benefit from this residual informational value. Sudden changes, in contrast, make past experiences immediately uninformative for future decisions. In these cases, relying on distant outcomes would be misleading. Because individuals rely more heavily than groups on recent experiences, they can be expected to show a greater advantage over groups in environments that change suddenly than in environments that change gradually.

#### **Expected patterns of adaptation**

Taken together, the evidence on group and individual decision making suggests that—due to their memory advantage—groups will behave more in accordance with Bayesian updating than individuals. Less reliable memory may lead individuals to perform better than groups after a change in the environment, showing a behavior that is more distant from the Bayesian benchmark. This pattern of relative adaptation should be more marked when change in the environment occurs suddenly.

### **The Present Experiment**

Our experiment consisted of six decision games. In each game, participants were asked to make 100 choices between two options either individually or as part of a group. In both options, participants could win or lose points with some probability. The number of points for each option remained constant throughout the 100 trials, but the probabilities of winning or losing could change. Each of the six games consisted of a nonstationary option (NS), where the probabilities changed after a certain number of trials, and a stationary option (S), where the probabilities did not change. Participants had no previous information about the options' outcomes or probabilities, and learned solely from feedback. After each choice, participants received feedback about the chosen option and the nonchosen option (i.e., the foregone payoff). Participants were not informed about the change in the underlying probabilities of the options, and could only perceive such change through observing the outcomes of their choices. The rationale for providing foregone payoffs is that we were interested in the adaptation of choice behavior to changes in the environment, and participants may otherwise have missed the opportunity to observe a change and react to it. Without notification of foregone payoffs, decision makers tend to avoid options that produce early unfavorable payoffs (Denrell & March, 2001). When early “hot stoves” occur, negative impressions tend to remain uncorrected.

Participants were randomly assigned to one of two conditions: They played either individually or as members of a three-person group. Participants in groups had complete freedom to decide how the group worked, but they were encouraged to collaborate in the group's decisions such that every choice was a collective decision.

## **Method**

### **Participants**


Eighty paid volunteers (65% male) from Carnegie Mellon University participated in the study. Twenty participants were assigned to the *individual* condition, and 60 participants were assigned to the *group* condition, forming 20 groups of three participants each.

### **Task and Apparatus**

Participants played six computer games, each consisting of 100 choices between two “money machines,” presented as two unlabeled buttons on a computer screen. In each trial, participants clicked on one of the two buttons and either won or lost points. After clicking, the participants saw the outcomes of their choices. They also saw the forgone payoffs (e.g., the payoffs they would have obtained had they chosen otherwise). Their objective was to maximize the number of points obtained over the 100 trials in each game. Each of the money machines represented a lottery with preprogrammed probabilities and payoffs. For the S option, the probability of winning was constant over the 100 repetitions. For the NS option, the probability of winning changed.

Figure 1 shows the payoff and probability structure of the six games. The games followed the structure of Rakow and Miler (2009), but we implemented a different compensation scheme. In games 1, 2, and 3, the probabilities of the NS option were programmed to change gradually. In games 4, 5, and 6, the probabilities of the NS option were programmed to change suddenly. In all

trials, one option was objectively better than (dominated) the other. All six games involved two stages: A stage where one option dominated the other, and a stage where the relationship between the options was reversed. We refer to the trial in which one option changes from better to worse as the reversal point. For example, in game 1, the NS option offered +10 points with .9 probability and −20 points otherwise in trials 1 to 20 (i.e., an expected value, EV, of 7). In trials 21 to 60, the probability of gaining +10 dropped by .01 per trial, so that by trial 61, the NS option offered +10 with .5 probability and −20 otherwise (i.e., EV of −5). The S option offered +10 with .7 probability and −20 otherwise across all trials (i.e., EV of 1). Thus, the EV of the NS option was higher than that of the S option in trials 1 to 40, but lower in trials 41 to 100.



*Figure 1.* Games used in the experiment. Each graph shows the probability of obtaining a gain in the nonstationary option (gray) and the stationary option (black) across the 100 trials. The amounts offered in each game are displayed inside the chart area of each graph. For example, game 1 is a choice between a nonstationary option that offers +10 or −20 and a stationary option that offers the same amounts. In the NS option, the probability of +10 is .9 from trials 1 to 20 and

then decreases by .01 per trial for 40 trials. From trial 61 to 100, the probability of +10 is .5. In the S option, the probability of +10 is .7 across all trials.

### **Design**

Because there was no precedent for the size of the effect studied in this experiment, we recruited 80 participants, matching the sample size in Rakow's and Miler's (2009) experiment 1. Participants were randomly assigned to either the *individual* ( $N = 20$ ) or the *group* ( $N = 60$ , thus 20 groups) condition. They played all six games described above in a random order. We also randomized the location of the NS option on the screen. Each game involved a stage where  $EV_{NS} > EV_S$  and a stage where the inequality was reversed.

### **Procedure**

Participants in the individual condition performed the experiment in individual computer booths. Participants in the group condition performed the task in an open space in single-group sessions. Groups worked on one computer on a table that allowed sufficient space for the three participants to collaborate comfortably. In both conditions, an experimenter read the following instructions (slightly modified from Rakow & Miler, 2009), which were identical for both conditions except for the variations shown in brackets:

You will play a series of games. In each game you will see two “money machines” like those shown below. When you click on a machine you will win or lose points. Your target [The target of your group] is to obtain as many points as possible.

In each game you will have 100 “goes” to win points. On each go, you [your group] should decide which machine you want to get points from. On each go, the probability of winning or losing points is fixed in advance. However, this probability may

change gradually from one go to the next.<sup>4</sup> When you click on a machine, you will see how many points you obtained. You will also see how many points you would have obtained from the other machine. [We encourage you to collaborate in the group's decisions such that every choice is a collective decision.]

If you have any questions, ask the experimenter. Press “Ready to begin” to start the first game.

The instructions warned participants that a change in probabilities *could* occur, although they did not say how prevalent the change would be—whether it would occur in some or all games—and also did not announce *when* in the sequence of trials the change would occur. Our goal was that participants would not assume a static environment.

All payoffs were compensated with money, with points being converted to cents. In the group condition, all three members received the full compensation (i.e., group members had the same economic incentives as individual participants). The mean individual compensation was U.S.\$7.90.

## Results

### How Did Groups and Individuals Adapt Their Choices After a Change?

Figure 2 plots the average individual and group rate of choices of the better option across trials. The general pattern is that groups made better choices than individuals before the reversal point (depicted by the vertical dashed line), but that individuals tended to choose the better option more often than groups after the reversal point. This pattern is evident in games 1 through 5, but less evident in game 6.

---

<sup>4</sup> Although the instructions indicated that probabilities “may change gradually,” in three games probabilities change suddenly. This was a mistake in the wording of our instructions, but we believe that they are sufficiently vague and that the task is sufficiently uncertain for this wording to have no qualitative impact on behavior.






Figure 2. Rate of choices of the better option. The vertical dashed lines indicate the point at which the dominance of one option over the other is reversed, i.e., the reversal point. The rate of choices is smoothed within each stage.

We performed a multilevel general linear model, with two between-subjects conditions (individuals and groups), six games, and two stages (before and after the reversal point) as within-subjects factors. The dependent measure was the rate of choices of the better option. The difference in adaptation between groups and individuals is described by the significant stage  $\times$  condition interaction,  $\chi^2(2) = 127.8$ ,  $p < .001$ ,  $r = .26$ , 95%  $CI = [.02, .06]$ . Before the reversal point, groups performed better ( $Mean = .77$ , 95%  $CI = [.73, .81]$ ) than individuals ( $Mean = .69$ , 95%  $CI = [.65, .73]$ ),  $p = .04$ . After the reversal point, however, this pattern was reversed,


with individuals performing better ( $Mean = .53$ , 95% CI = [.49, .58]) than groups ( $Mean = .45$ , 95% CI = [.40, .50]),  $p = .10$ . Both post hoc tests were Bonferroni corrected.

There was no significant main effect of condition,  $\chi^2(1) < .01$ ,  $p = .99$ ,  $r < .01$ , 95% CI = [−.03, .03], indicating that the rate of choices of the better option was similar for groups and individuals across the two stages and the six games. The stage of the game had a significant main effect, however,  $\chi^2(6) = 116.4$ ,  $p < .001$ ,  $r = .62$ , 95% CI = [−.14, −.10]. Both groups and individuals made better decisions before the reversal point than after the reversal point, indicating that adaptation was difficult for both individuals and groups.

The difference in adaptation for groups and individuals was not sensitive to whether the change was sudden or gradual. A comparison of optimal choices after gradual change (games 1–3) versus sudden change (games 4–6) suggests that individuals did not adapt better in sudden ( $Mean = .54$ , 95% CI = [.48, .60]) than in gradual games ( $Mean = .53$ , 95% CI = [.47, .59]),  $p = 1$ . Groups also showed a similar level of adaptation in both sudden ( $Mean = .46$ , 95% CI = [.38, .53]) and gradual games ( $Mean = .45$ , 95% CI = [.38, .52]),  $p = 1$ . Both post hoc tests were Bonferroni corrected.

**How did groups and individuals adapt relative to the Bayesian benchmark?** The Bayesian updating model was applied to the sequences of outcomes observed by each individual and group. Figure 3 shows the rate of choices of the nonstationary option across trials, for both individuals and groups, relative to their corresponding Bayesian benchmark. As the figure suggests, groups behaved more in accordance with the Bayesian benchmark than individuals did—the dark thin and thickest curves are closer to each other than are the light thin and thickest curves across problems. We calculated the percentage of choices that corresponded to the Bayesian benchmark. Before the reversal point, 74% of group choices corresponded to Bayesian

updating, relative to 67% of individual choices. This pattern persisted after the reversal point, with 68% and 62% for groups and individuals, respectively. The overall pattern suggests that, on average, both individuals and groups perceived a change and discounted earlier experiences beyond the Bayesian benchmark. This pattern was more pronounced in individuals than in groups, however.



*Figure 3.* Rate of choices of the nonstationary option. The vertical dashed lines indicate the reversal point. Three types of curve appear in each graph. (1) Thin lines indicate mean observed behavior; (2) thicker lines indicate the smoothed observed behavior; and (3) the thickest lines indicate the mean predictions of the Bayesian updating model applied to each individual's and group's sequences of outcomes.

Which factors underlie this pattern of results? First, groups made better choices than individuals before the reversal point, a result that is consistent with groups being better information processors. After the reversal point, however, groups did not adapt as much as individuals, and behaved more in line with a Bayesian benchmark that assumes perfect memory. Next, we examine the role of memory and of routines as possible drivers of this pattern of results.

### **Is the Behavior of Groups Consistent with Them Having Better Memory?**


We studied memory in individual and group decisions by modeling the respective behavior using an Instance-Based Learning (IBL) model for repeated binary choice (Lejarraga et al., 2012) that originated from IBL theory (IBLT; Gonzalez et al., 2003). In previous work (Lejarraga et al., 2012), we used this IBL model to predict the average individual behavior reported in Rakow and Miler (2009). Although the model was not fitted to the individual's choices but to mean choices, predictions from the IBL model provided an accurate account of the average behavior relative to other, comparable models. The IBL model assumes that individuals learn by activating memories about past experiences. It thus learns from observed outcomes (both obtained and foregone) and makes trial-by-trial predictions. Outcomes that have been observed more often or more recently will be more active in memory than outcomes that are infrequent or distant. The activation of past experiences in the IBL model is influenced by two assumed cognitive processes: memory decay and retrieval noise. Because memories decay with time (or trials), old memories are less active in memory than recent ones (i.e., recency). Decay is modulated by a free parameter  $d$  that is a negative power of the time elapsed since a given outcome was experienced. Therefore, memories decay exponentially (an assumption derived from the ACT-R theory of cognition, Anderson & Lebiere, 1998). A lower decay parameter

signifies better memory. Therefore, lower decay values make the IBL model behave more in line with the predictions of the Bayesian updating model. When fitting the two parameters of this IBL model to individual and group choices, we expect to find a lower decay parameter in groups than in individuals.

The IBL model also assumes that memory retrieval is not perfect, but influenced by random noise. The probability of retrieval of each outcome is given by the relative activation of one outcome with respect to the other outcomes. The relative activation is altered by a noise parameter  $\sigma$ . With higher noise, the probability of retrieval of two experienced outcomes will be more equal, whereas with low noise, memory retrieval is more accurate. We expect that retrieval noise is lower in groups than in individuals. Note that IBL is not a model of group choice; rather, we use this IBL model as a tool to interpret whether group choices were made “as if” the decision maker had a better memory.

We fitted this IBL model to the decisions of the 20 individuals and 20 groups in the experiment (see Appendix for a full model specification and fitting procedure). The resulting memory decay and noise parameter values were indeed lower for groups than for individuals. Model performance is displayed in Figures 4 and 5, for individuals and groups respectively. Figure 6 shows the goodness of fit ( $G^2$ ) for the IBL model’s choices as a function of different sets of parameters  $d$  and  $\sigma$ . Note that the IBL model produced qualitatively better predictions of better memory (lower  $d$  and  $\sigma$ ) for groups than for individuals, a result that is consistent with groups behaving more in line with the predictions of the Bayesian updating model. These results are also consistent with previous evidence that groups have better memory than individuals, and that better memory can be detrimental to adaptation in changing environments. However, these results may also be the consequence of groups adopting routines. As discussed above, groups

may adopt routines to evade the costs of coordinating a response, but these routines also limit group adaptation. We consider this possibility next.



*Figure 4.* Rate of choices of the nonstationary option across trials for individuals and for the IBL model fitted to each individual's choices.




Figure 5. Rate of choices of the nonstationary option across trials for groups and for the IBL model fitted to each group's choices.




Figure 6. Models goodness of fit ( $G^2$ ) for group and individual choices. The dark point indicates the best-fitting parameter set.

**Is the Behavior of Groups Consistent with the Adoption of Routine Strategies?**


The groups' limited adaptation may have been caused by their adoption of a routine strategy. To address this issue, we examined two well-studied routine strategies: majority and unanimity (e.g., Tindale & Kameda, 2000). These two rules for aggregating individual preferences are routines to the extent that they provide the group with a solution to avoid engaging in costly negotiations in which individual members try to persuade one another. If the observed group behavior can be reproduced by aggregating simulated individual preferences in accordance with these routines, then we need to de-emphasize the effect of memory. However, if the choices produced by aggregating individual behavior are qualitatively different from the observed group behavior, then we can interpret the result as additional support for the memory explanation.

We modeled group behavior by creating three simulated participants using the IBL model with the best-fitting parameters for individual behavior ( $\sigma = 0.6$ ,  $d = 0.9$ ). The simulation consisted of having these three simulated participants make choices individually given the sequences of outcomes they observed in the six games. To generate predictions by majority, we combined the individual choices for each of the three simulated participants according to a simple majority rule. In each trial, a group made an NS choice if at least two of its members made an NS choice. To generate predictions by unanimity, we proceeded as follows: A group made decisions by majority in the first trial and subsequent decisions were altered only if there was coincidence among all three group members.

Results are displayed in Figure 7. Visual inspection suggests that the majority and unanimity routines captured the general tendency of group behavior but—with the exception of



game 6—overrepresented the ability of groups to adapt after the change. In games 1 through 5, the simulated routine strategies adapted earlier and to a greater extent than the real groups did. In general, routine strategies produced behavior that was more in line with the better choice than was actual group behavior. These results do not rule out routines as a possible cause of group choice, but they do not provide convincing evidence that routine strategies are the main factor driving group behavior in decisions from experience.



*Figure 7.* Rate of choices of the better option by groups and by majority and unanimity routines.

The curves were smoothed for clarity.


### **Did Individuals React Better or Were They Favored by Change in the Environment?**

Is it possible that individuals did not react better than groups, but that they fared better after change in the environment because they were initially less committed to the initially better

option? One way to address whether individuals were less committed than groups is by looking at how frequently they alternated between options (e.g., choosing S in one trial and NS in the following trial counts as one alternation). Frequent alternations suggest lower commitment to a particular option. In our experiment, participants received feedback on both the options they chose and those they did not choose, meaning that they did not need to switch options to obtain new information. Therefore, the frequency of alternations was not confounded with participants' tendency to engage in explorative behavior (Hertwig et al., 2009).

We calculated the rate of alternations across individuals and groups for each trial (Figure 8). Results of a multilevel general linear model indicated that individuals alternated more ( $Mean = .26$ , 95% CI = [.23, .28]) than did groups ( $Mean = .19$ , 95% CI = [.16, .21]),  $\chi^2(1) = 3.73$ ,  $p = .05$ ,  $r = .30$ , 95% CI = [-.002, .14]. These results suggest that individuals were less committed to their choices. Individual behavior, when compared with that of groups, was better prepared to changes in the environment. But can commitment exclusively explain the better performance of individuals after the reversal point?

Close examination of Figure 3 suggests that commitment is not the main driver. The individual and group curves run close to parallel in all problems, suggesting that groups and individuals changed their behavior to a similar extent. We calculated the amount of change in the rate of NS choices from the first to the second stage for groups and individuals in each problem. The mean change for groups and individuals amounted to 22% and 23%, respectively. This pattern indicates a similar level of change in individual and group decisions from before to after the reversal point.



*Figure 8.* Rate of alternations between options. The vertical dashed lines indicate the reversal point. The thin lines indicate mean observed behavior. The thicker lines indicate the smoothed observed behavior.

### Discussion and Conclusion

We were interested in the relative performance of groups and individuals after a change in the decision environment. We used a task in which the change was not announced; participants could detect it only from experiencing the outcomes of their choices.

Consistent with the previous literature, we found that groups performed better than the average individual while the decision task was stable. However, their performance was no longer superior after a change in the decision environment: The average individual now made better choices than did groups. How do we explain this behavior?

Coordination imposes costs to group decision making (Steiner, 1972). Groups may therefore adopt routine strategies to avoid the burden of a coordinated response, particularly in recurrent settings, as examined in this study. We simulated two plausible routine strategies that may allow groups to sidestep this costly coordination: majority and unanimity. The two simulated routines captured the general pattern, but deviated qualitatively from the observed group behavior: after the change, the simulated routines adapted sooner and to a greater extent than the real groups did. We interpret this result as suggesting that the groups did not use explicit routine strategies in this task. This observation is supported by the fact that groups alternated between options significantly less often than the individuals did, as if it were more costly for groups to coordinate a change in a selected course of action. Overall, we think that the evidence for the adoption of routine strategies as the main motor of group behavior is weak in this task.

Group behavior was consistent with groups having better memory than individuals. Their behavior was closer to a Bayesian benchmark, which assumes perfect memory and no change in the environment until sufficient evidence accumulates to indicate so. Moreover, we fitted a learning model (IBL) to group behavior, and the resulting parameters suggest that our groups behaved like individuals with longer and more accurate memory. Is it possible that coordination among group members produces behavior consistent with better memory? We believe that group discussions may allow group members to correct each other's memories, and may activate older memories that may have decayed in an individual setting. Error correction has been argued to improve joint memory (Hinsz, 1990; Vollrath et al., 1989). Nevertheless, the pattern of results cannot be explained solely by the effect of joint memory, as we found no differences in adaptation in games with sudden versus gradual changes. Our results also suggest that individuals were less committed to a particular option before the change in the environment—as

reflected by their higher alternation rates—and were therefore less exposed to detrimental effects of change.

Our results are consistent with Rakow's and Miler's (2009) finding that a memory aid impedes adaptation in a changing environment. In our study, however, the memory support was provided by group interaction—presumably by means of error correction—and not by an external aid.

### **Open Questions and Future Research**

Although behavior in our experiment was consistent across games 1 to 5, game 6 showed a divergent pattern of results. In contrast to the general pattern of groups making better choices before the reversal point, and individuals tending to choose the better option after it, there was no difference between groups and individuals in this game; if anything, they showed the opposite pattern. We have no explanation for this result. Game 6 was one of two games in which the NS option went from worse to better (the other was game 3). As discussed by Rakow and Miler (2009), it is reasonable to expect that more attention is paid to the option that dominates in early trials. Thus, it may be easier to detect a worsening NS option than an improving one. A group with more “eyes” tracking the environment (Krause & Ruxton, 2002) may be better able than an individual to spot an improving NS option, and thus catch up with the behavior of individuals after change. This argument does not explain the results of game 3, however. Whether groups are better able to adapt to less noticeable changes is an interesting question that should be addressed in future research.

We believe that the study of decisions from experience has potential to uncover important aspects of group and individual adaptation. In this experiment, individuals showed more conservative (i.e., less committed) behavior than groups, which was favorable in the

environments examined. Follow-up studies could examine how key properties of the environment influence adaptation (as per Kämmer et al., 2013). For example, a change in the environment could make the currently attractive option even more attractive—which would presumably favor groups—or changes could be cyclical. The environment could include mixed tasks, some in which payoff distributions change and some in which they do not, potentially leading individuals to overreact to random streaks of outcomes.

Equally important, follow-up studies could examine social dynamics within groups. For example, it is reasonable to expect that groups with leaders and groups without leaders may interact differently and therefore differ in the way they process information and adapt to changes in the environment. Group size—and, in particular, whether groups have an even or odd number of members—may also have an influence (Laughlin, Hatch, Silver, & Boh, 2006). Whether a majority rule is sufficient to make a decision may have implications for the rate of alternation and, subsequently, for adaptation. Other costs of coordination may also affect adaptation, such as the level of diversity within the group (Nijstad & Kaps, 2008).

Further research could examine how groups and individuals differ in their ability to detect a change and act accordingly. Experiments could address this question by including a measure of recognition that the environment has changed. Learning more about how groups and individuals differ in the processes that lead to adaptation has important implications for the practice of group behavior and could be relevant for the literature on organizational behavior.

Overall, our results suggest that an individual restaurateur will be better able than a group to adapt to changes in the quality of his suppliers, because although he may not always rely on the best supplier, he will pay more attention to the quality of recent deliveries, and will thus be quicker to drop those who have changed from better to worse. Groups responsible for making

recurrent decisions may benefit from the perspective of outside individuals, who do not share the memories of group members and may consequently be more sensitive to events of the recent past.

### References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16, 215–233.  
Doi: 10.1002/bdm.443
- Beersma, B., Hollenbeck, J. R., Conlon, D. E., Humphrey, S. E., Moon, H., & Ilgen, D. R. (2009). Role negotiation in self-managed teams: The effects of history and composition on coordination and performance. *Organizational Behavior and Human Decision Processes*, 108, 131–142. doi:10.1016/j.obhdp.2008.07.002
- Betsch, T. (2005). Preference theory: An affect-based approach to recurrent decision making. In T. Betsch & S. Haberstroh (Eds.). *The routines of decision making* (pp. 39–65). Mahwah, NJ: Lawrence Erlbaum.
- Betsch, T., Fiedler, K., & Brinkmann, J. (1998). Behavioral routines in decision making: The effects of novelty in task presentation and time pressure on routine maintenance and deviation. *European Journal of Social Psychology*, 28, 861–878 .  
doi:10.1002/(SICI)1099-0992(1998110)28:6<861::AID-EJSP899>3.0.CO;2-D
- Betsch, T., Haberstroh, S., Glöckner, A., Haar, T., & Fiedler, K. (2001). The effects of routine strength on information acquisition and adaptation in recurrent decision making. *Organizational Behavior and Human Decision Processes*, 84, 23–53.  
doi:10.1006/obhd.2000.2916



- Betsch, T., Haberstroh, S., & Höhle, C. (2002). Explaining and predicting routinized decision making: A review of theories. *Theory and Psychology, 12*, 453–488.  
doi:10.1177/0959354302012004294
- Betsch, T., Lindow, S., Engel, C., Ulshöfer, C., & Kleber, J. (in press). *Has the world changed? My neighbor might know effects of social context on routine deviation*. Retrieved from [http://www.coll.mpg.de/pdf\\_dat/2011\\_21online.pdf](http://www.coll.mpg.de/pdf_dat/2011_21online.pdf)
- Betts, K. R., & Hinsz, V. B. (2010). Collaborative group memory: Processes, performance, and techniques for improvement. *Social and Personality Psychology Compass, 4*, 119–130.  
doi:10.1111/j.1751-9004.2009.00252.x
- Biele, G., Rieskamp, J. and Gonzalez, R. (2009), Computational Models for the Combination of Advice and Individual Learning. *Cognitive Science, 33*, 206–242. doi: 10.1111/j.1551-6709.2009.01010.x
- Bröder, A., & Schiffer, S. (2006). Adaptive flexibility and maladaptive routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 32*, 904–918. doi:10.1037/0278-7393.32.4.904
- Charness, G., & Sutter, M. (2012). Groups make better self-interested decisions. *Journal of Economic Perspectives, 26*, 157–176. doi:10.1257/jep.26.3.157
- Cohen, T., & Thompson, L. (2011). When are teams an asset in negotiations and when are they a liability? In E. Mannix, M. Neale, & J. Overbeck (Eds.), *Research on managing groups and teams: Negotiation in groups* (Vol. 14, pp. 3–34). Bingley, UK: Emerald.
- Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science, 12*, 523–538. doi:10.1287/orsc.12.5.523.10092

- Erev, I., Ert, E. & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, 21, 575–597.  
doi: 10.1002/bdm.602
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., Hertwig, R., Stewart, T., West, R. and Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23, 15–47.  
doi: 10.1002/bdm.683
- Gersick, C. J., & Hackman, R. J. (1990). Habitual routines in task-performing groups. *Organizational behavior and human decision processes*, 47, 65–97. doi:10.1016/0749-5978(90)90047-D
- Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, 27, 591–635. doi:10.1207/s15516709cog2704\_2
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15, 534–539.  
doi:10.1111/j.0956-7976.2004.00715.x
- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, 13, 517–523. doi:10.1016/j.tics.2009.09.004
- Hill, G. W. (1982). Group versus individual performance: Are  $n + 1$  heads better than 1? *Psychological Bulletin*, 91, 517–539. doi:10.1037/0033-2909.91.3.517
- Hinsz, V. B. (1990). Cognitive and consensus processes in group recognition memory performance. *Journal of Personality and Social Psychology*, 59, 705–718.  
doi:10.1037/0022-3514.59.4.705

- Hinsz, V. B., Tindale, R. S., & Vollrath, D. A. (1997). The emerging conceptualization of groups as information processors. *Psychological Bulletin*, 121, 43–64. doi:10.1037/0033-2909.121.1.43
- Hollenbeck, J. R., Ellis, A. P. J., Humphrey, S. E., Garza, A. S., & Ilgen, D. R. (2011). Asymmetry in structural adaptation: The differential impact of centralizing versus decentralizing team decision-making structures. *Organizational Behavior and Human Decision Processes*, 114, 64–74. doi:10.1016/j.obhdp.2010.08.003
- Johnson M. D., Hollenbeck J. R., Humphrey S. E., Ilgen D. R., Jundt D., & Meyer C. J. (2006) Cutthroat cooperation: Asymmetrical adaptation to changes in team reward structures. *Academy of Management Journal*, 49, 103–119. doi:10.5465/AMJ.2006.20785533
- Kämmer, J. E., Gaissmaier, W., & Czienskowski, U. (2013). The environment matters: Comparing individuals and dyads in their adaptive use of decision strategies. *Judgment and Decision Making*, 8, 299–329.
- Katz, L. (1964). Effects of differential monetary gain and loss on sequential two-choice behavior. *Journal of Experimental Psychology*, 68, 245–249. doi:10.1037/h0044150
- Krause, J., & Ruxton, G. D. (2002). *Living in groups*. Oxford, UK: Oxford University Press
- Laughlin, P. R., Hatch, E. C., Silver, J. S., & Boh, L. (2006). Groups perform better than the best individuals on letters-to-numbers problems: Effects of group size. *Journal of Personality and Social Psychology*, 90, 644–651. doi:10.1037/0022-3514.90.4.644
- Lejarraga T., Dutt V., & Gonzalez C. (2012). Instance-based learning: A general model of repeated binary choice. *Journal of Behavioral Decision Making*, 25, 143–153. doi:10.1002/bdm.722

- LePine, J. A. (2003) Team adaptation and postchange performance: Effects of team composition in terms of members' cognitive ability and personality. *Journal of Applied Psychology*, 88, 27–39. doi:10.1037/0021-9010.88.1.27
- LePine, J. A. (2005) Adaptation of teams in response to unforeseen change: Effects of goal difficulty and team composition in terms of cognitive ability and goal orientation. *Journal of Applied Psychology*, 90, 1153–1167. doi:10.1037/0021-9010.90.6.1153
- Marks, M. A., Zaccaro, S. J., & Mathieu, J. E. (2000). Performance implications of leader briefings and team-interaction training for team adaptation to novel environments. *Journal of Applied Psychology*, 85, 971–986. doi:10.1037/0021-9010.85.6.971
- Myers, J. L., & Sadler, E. (1960). Effects of range of payoffs as a variable in risk taking. *Journal of Experimental Psychology*, 60, 306–309. doi:10.1037/h0042499
- Nijstad, B. A., & Kaps, S. C. (2008). Taking the easy way out: Preference diversity, decision strategies, and decision refusal in groups. *Journal of Personality and Social Psychology*, 94, 860–870. doi:10.1037/0022-3514.94.5.860
- Olsson, A-C., Juslin, P., & Olsson, H. (2006). Multiple cue judgment in individual and dyadic learning. *Journal of Experimental Social Psychology*, 42, 40–56.  
doi:10.1016/j.jesp.2005.01.004
- Rakow, T., & Miler, K. (2009). Doomed to repeat the successes of the past: History is best forgotten for repeated choices with nonstationary payoffs. *Memory and Cognition*, 37, 985–1000. doi:10.3758/MC.37.7.985
- Reimer, T., Bornstein, A.L., & Opwis, K. (2005). Positive and negative transfer effects in groups. In T. Betsch & S. Haberstroh (Eds.), *The routine of decision making* (pp. 175–192). Mahwah, NJ: Lawrence Erlbaum Associates.

- Ritter, B. A., & Lord, R. G. (2007). The impact of previous leaders on the evaluation of new leaders: An alternative to prototype matching. *Journal of Applied Psychology, 92*, 1683–1695. doi:10.1037/0021-9010.92.6.1683
- Schneider, W., & Shiffrin, R. M. (1997) Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84*, 1–66. doi:10.1037/0033-295X.84.1.1
- Steiner, I. D. (1972). *Group process and productivity*. New York, NY: Academic Press.
- The world's 50 best restaurants. (2013). Retrieved August 1, 2013, from <http://www.theworlds50best.com/list/1-50-winners/>
- Tindale, R. S., & Kameda, T. (2000). 'Social sharedness' as a unifying theme for information processing in groups. *Group Processes & Intergroup Relations, 3*, 123–123. doi:10.1177/1368430200003002002
- Vollrath, D. A., Sheppard, B. H., Hinsz, V. B., & Davis, J. H. (1989). Memory performance by decision-making groups and individuals. *Organizational Behavior and Human Decision Processes, 43*, 289–300. doi:10.1016/0749-5978(89)90040-X
- Yechiam, E. & Busemeyer, J. R. (2005). Comparisons of basic assumptions embedded in learning models for experienced based decision making. *Psychonomic Bulletin and Review, 12*, 387–402. doi:10.1016/j.geb.2007.08.011

## Appendix

### Instance-Based Learning Model

The Instance-Based Learning (IBL) model (Gonzalez et al., 2003; Lejarraga et al., 2012) evaluates options according to their blended value. The *blended value*  $V$  of option  $j$  is

$$V_j = \sum_{i=1}^n p_i x_i,$$

where  $x_i$  is the value of the observed outcome  $i$  and  $p_i$  is the probability of retrieval of that outcome from memory. At trial  $t$ , the *probability of retrieval* of observed outcome  $i$  is a function of the activation of that outcome relative to the activation of all the observed outcomes  $k$  in  $j$ :

$$P_{i,t} = \frac{e^{A_{i,t}/\tau}}{\sum_i^k e^{A_{k,t}/\tau}},$$

where  $\tau$  is random noise defined as  $\tau = \sigma\sqrt{2}$ , and  $\sigma$  is a parameter fitted to the data (see below). At trial  $t$ , the activation (Anderson & Lebiere, 1998) of an outcome  $i$  is

$$A_{i,t} = \sigma \ln \left( \frac{1-\gamma_{i,t}}{\gamma_{i,t}} \right) + \ln \sum_{t_p \in \{1, \dots, t-1\}} (t - t_p)^{-d},$$

where  $d$  is a decay parameter fitted to the data (see below);  $\gamma_{i,t}$  is a random draw from a uniform distribution bounded between 0 and 1 for each outcome and trial; and  $t_p$  is each of the previous trial indexes in which the outcome  $i$  was encountered.

The IBL model chooses the option with the highest blended value  $V_{j,t}$ .

### Estimating Parameters for the IBL Model

The IBL model assumes a deterministic choice rule, but involves a random component in the activation function modulated by  $\gamma_{i,t}$ . The model therefore predicts different choices depending on the particular realization of  $\gamma_{i,t}$ . To control for this stochasticity, we generated 1,000 predictions for each participant (whether an individual or a group) and game. Therefore, for each participant and game, the probability of the model making a correct prediction was the proportion of the 1,000 correct predictions. We then calculated the model's fit according to its discrepancy from the observed choices as measured by the negative log likelihood:

$$G^2 = -2 \sum_{i=1}^n \ln[f_i(y|\theta)],$$

where  $n$  is the pooled number of choices ( $20 \text{ participants} \times 6 \text{ games} \times 100 \text{ trials}$ ), and  $f_i(y|\theta)$  is the probability with which IBL makes a correct prediction based on a particular parameter set  $\theta$ . Each particular set of parameters consisted of a value of  $d$  and a value of  $\sigma$ . We tested all parameter combinations with 0.1 increments for  $0 < d < 6$  and  $0 < \sigma < 3$ ; results are reported in Figure 6.